

CERTIFICATE OF MAILING UNDER 37 CFR§ 1.10

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail in an envelope addressed to: Assistant Commissioner of Patents, Washington, DC 20231 on January 24, 2002

EXPRESS MAIL LABEL: EL 888550449 US

Amirah Scarborough
Name of Person Mailing Document

Amirah Scarborough
Signature of Person Mailing Document

LOGGING INSERTION/REMOVAL OF SERVER BLADES
IN A DATA PROCESSING SYSTEM**BACKGROUND****1. Field of the Present Invention**

The present invention generally relates to the field of data processing systems and more particularly to a system and method for logging the insertion and removal of circuit boards or blades in a data processing system.

2. History of Related Art

Data processing systems in general and server-class systems in particular are frequently implemented with a server chassis or cabinet having a plurality of racks. Each cabinet rack can hold a rack mounted device (also referred to herein as a blade or server blade) on which one or more general purpose processors and/or memory devices are attached. The racks are vertically spaced within the cabinet according to an industry standard displacement (the "U"). Cabinets and racks are characterized in terms of this dimension such that, for example, a 42U cabinet is capable of receiving 42 1U rack-mounted devices, 21 2U devices, and so forth. Dense server designs are also becoming available, which allow a server chassis to be inserted into a cabinet

rack, thus allowing greater densities than one server per 1U. To achieve these greater densities, the server chassis may provide shared components, such as power supplies, fans, or media access devices which can be shared among all of the blades in the server blade chassis.

In a rack-mounted server blade environment, the ability to hot plug server blades is a standard feature. Hot plugging refers to the ability to install and remove a blade without turning off power to the cabinet. When a new server blade is initially installed into a cabinet, the blade generally contains no operating system and no persistent data. Making a newly added blade functional requires deployment software that is capable of recognizing that a new blade has been added, determining the blade characteristics to uniquely identify the blade, powering the blade on, and assigning a functional boot image to the blade. For purposes of this disclosure, a boot image refers generally to software stored in persistent storage that is executed following a power-on or system reset event. The boot image may execute a self test (commonly referred to as a power on self test or POST), load a basic I/O system (BIOS) into memory, and install a functional operating system.

While the rack design beneficially facilitates the easy configuration and expansion of server systems, it also allows server blades to be moved about freely within a cabinet or among different cabinets. In this environment, it is frequently difficult for system administrators to keep track of the location of individual server blades and how many times a server blade has been inserted into a rack. For a variety of reasons, it is highly desirable to maintain accurate information about the location of each server blade. Server blades are typically warranted by their manufacturer for a specified period of time that typically begins when the server blade is first placed in service. Determining the warranty periods of individual blades can quickly prove formidable if the location of each blade is not carefully maintained. Similarly, it is desirable to maintain service logs for each server blade to identify problematic blades and determine the overall cost of operating the blades. Moreover, server blades are typically specified for some maximum of insertions. If blade movement within a system is not tracked, a server blade may be inserted more times than permitted by the blade specification thereby potentially resulting in faulty operation and/or the denial of warranty benefits by the manufacturer. For at least these reasons, it is desirable to implement a system and method in which server blade insertion and movement logs are kept.

SUMMARY OF THE INVENTION

The problems identified above are in large part addressed by a data processing system, referred to as a blade, comprising at least one main processor connected to a system bus, a system memory connected to the system bus and accessible to each of the main processors, a tamper mechanism, and a local service processor. The tamper mechanism is configured to change state each time the system is inserted into a slot in a rack enclosure. The local service processor on each blade is connected to the tamper mechanism and configured to update an insertion log upon detecting a change in state of the tamper mechanism. The insertion log provides a history of at least some rack insertions to which the system has been subjected. The system may include a non-volatile storage element accessible exclusively to the local service processor that contains the insertion log. The insertion log may include an insertion counter. In this embodiment, the local service processor is configured to increment the insertion counter upon each insertion. The local service processor may be further configured to issue an alert if the insertion counter exceeds a predetermined value. In one embodiment, the system further includes a battery backed real-time clock connected to the local service processor. The local service processor is configured to include real-time information corresponding to each insertion event in the insertion log. Each entry in the insertion log may include the identity of the rack enclosure and the geographical address of the slot of the corresponding insertion event. The local service processor may be configured to detect the tamper mechanism state and update the insertion following a power event such that the insertion log update is independent of configuring the data processing system with a boot image.

BRIEF DESCRIPTION OF THE DRAWINGS

Other objects and advantages of the invention will become apparent upon reading the following detailed description and upon reference to the accompanying drawings in which:

FIG 1 is a block diagram of selected elements of a data processing system suitable for use with one embodiment of the present invention;

FIG 2 depicts a data processing network according to one embodiment of the invention;

FIG 3 illustrates additional detail of the data processing network of FIG 2;

FIG 4 illustrates selected details of a local service processor and a corresponding non-volatile log storage device according to one embodiment of the invention; and

FIG 5 illustrates additional detail of the log storage device of FIG 4 according to one embodiment of the invention.

5 While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description presented herein are not intended to limit the invention to the particular embodiment disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within
10 the spirit and scope of the present invention as defined by the appended claims.

DETAILED DESCRIPTION OF THE INVENTION

Turning now to the drawings, FIG 1 is a block diagram of selected features of a data processing system 100 suitable for implementing the present invention. Data processing system 100 may be implemented entirely upon a single printed circuit board. In this embodiment, data processing system 100 may be alternatively referred to herein as server blade 100. In the depicted embodiment, server blade 100 includes a set of main processors 102A through 102N (generically or collectively referred to as processor(s) 102) that are connected to a system bus 104. A common system memory 106 is accessible to each processor 102 via system bus 104. The system memory is typically implemented with a volatile storage medium such as an array of dynamic random access memory (DRAM) devices. The depicted architecture of server blade 100 is frequently referred to as a symmetric multiprocessor (SMP) system because each processor 102 has substantially equal access to system memory 106.

25 In server blade 100, a bus bridge 108 provides an interface between system bus 104 and an I/O bus 110 to which one or more peripheral devices 114A through 114N (generically or collectively referred to as peripheral device(s) 114) as well as a general purpose I/O (GPIO) port are connected. Peripheral devices 114 may include devices such as a graphics adapter, a high-speed network adapter or network interface card (NIC), a hard-disk controller, and the like. I/O
30 bus 110 is typically compliant with one of several industry standard I/O bus specifications

including, as a common example, the Peripheral Components Interface (PCI) bus as specified in *PCI Local Bus Specification Rev 2.2* by the PCI Special Interest Group (www.pcisig.com).

The depicted embodiment of server blade 100 includes a local service processor 116 connected to GPIO port 112. Local service processor 116 is configured to provide support for main processors 102. This support may include, for example, monitoring the power supplied to main processor(s) 102 and, in the event of a blade crash, initiating a restart of the main processors.

Turning now to FIG 2, a data processing network 200 according to one implementation of the present invention is depicted. Data processing network 200 includes a rack enclosure or cabinet 201 that includes a plurality of slots or racks 202. Each rack 202 is configured to receive a rack mounted device such as a server blade chassis identified by reference numerals 101a through 101n (generically or collectively referred to as server blade chassis 101) via a suitable connection mechanism such as a traditional edge connector. Each server blade chassis 101 typically contains one or more server blades 100 as described with respect to FIG 1. In one implementation, each server blade chassis 101 is a 4U component that may include as many as 16 server blades 100. Thus, the depicted embodiment of network 200 includes a set of server blade chassis 101, each of which includes one or more server blades 100. Each server blade 100 within server blade chassis 101a through 101n is typically connected to a local area network (LAN) through its NIC(s) such that server blades 100 in system 200 can communicate with each other. Network 200 may itself be connected to an external network such as the Internet through a gateway (not depicted) or other suitable network device.

Portions of the present invention may be implemented as a sequence of processor executable instructions (software) for locally maintaining a log of the insertion and removal of a server blade in a data processing network where the instructions are stored on a computer readable medium. During execution, portions of the software may reside in a volatile storage element such as the system memory 106 depicted in FIG 1 or an external or internal cache memory (not depicted) of main processor(s) 102. At other times, portions of the software may be stored on a non-volatile storage medium such as a floppy diskette, hard disk, CD ROM, DVD, magnetic tape, or other suitable storage medium. In addition, portions of the software may be executed by management module 120 while other portions are executed by local service processors 116 of each server blade 100.

Turning now to FIG 3, the local service processor 116 of each server blade 100 is connected to a set of physical identification (PID) connectors 130 that typically comprise a portion of the edge connectors of server blade 100. PID connectors 130 are configured to connect to a corresponding set of PID pins 134 of an edge connector 132 residing in each slot 202 of cabinet 201. The PID pins 134 of each connector 132 are uniquely wired to either ground or VDD to produce a binary value indicative of the physical position of the slot 202 in which the connector is positioned. The PID connectors 130 thus provide local processor 116 with means for determining the physical slot or geographical address occupied by the corresponding server blade 100.

Server blades 100 as depicted in FIG 3 are further configured with a tamper latch 135. Tamper latch 135 comprises a mechanism that indicates whether a corresponding device or apparatus has been altered. In the case of server blades 100, tamper latch 135 may be configured to transition from a first state to a second state when the corresponding server blade is inserted into (or removed from) its slot. Tamper latch 135 may employ an impedance element that is electrically altered when the corresponding server blade 100 is removed from its slot. The impedance element preferably remains in this altered state until reset by system software. Thus, tamper latch 135 may include a persistent impedance element that changes impedance when its corresponding server blade 100 is inserted into or removed from its slot.

The local service processors 116 of server blades 100 are directly connected to system power connector pins of their server blades such that local service processors 116 are powered on whenever they are installed in their slots 202. In response to a power-up event (meaning a power-on of the local service processor or a hardware reset), local service processors 116 immediately read their corresponding geographical address or slot location as well as chassis identification information (that is provided to them through separate connectors (not depicted)). In addition, a local service processor will read its corresponding tamper latch to determine if the blade has been removed since the last power up event. If the tamper latch is broken (in an impedance condition indicating that the blade has been removed), local service processor 116 increments the insertion count, then resets tamper latch 135. In addition, local service processor 116 is configured to log insertion information in a non-volatile insertion log.

Referring now to FIG 4 and FIG 5, one embodiment of system 100 has local service processor 116 connected to a non-volatile storage element 142. Storage element 142 is preferably writeable only by local server processor 116 (i.e., is read-only to main processors 102) and may be implemented with a flash memory card or other suitable non-volatile storage device.

5 Storage element 142 includes information comprising a blade insertion log 144.

Blade insertion log 144 typically includes a detailed history of the movement of server blade 100 since its original in-service date. Blade insertion log 144 may include, for example, a blade insertion table 152. Table 152 as depicted in FIG 5 includes a row for each insertion event. Each row of table 152 indicates the time of the corresponding insertion event, identification
10 information about the chassis in which the blade was inserted, as well as the physical slot location (geographical address) within the chassis.

In addition, blade insertion log 144 may include an insertion counter 150. If, in response to a power on event or hardware reset, local service processor 116 determines that its tamper latch 135 has been broken, the service processor will increment insertion counter 150 and update blade insertion log 144 by including a new row of information in the log. If the number of rows in log 144 is limited, local service processor may be configured to overwrite the oldest insertion data when the log is full. Optionally, the local service processor may be configured to only log insertion events where the physical position of the blade indicates that the blade has been
15 relocated since the last insertion.

The insertion counter 150 provides a secure count of the number of physical insertion to which the server blade has been subjected. If the insertion counter 150 exceeds a predetermined value, local service processor 116 may issue some form of alert, warning message, interrupt, or the like to indicate that the blade has been inserted/remove more times than permitted by specification thereby alerting the system administrator to a potential problem. The depicted
20 embodiment of system 100 further includes a real-time clock (RTC) 146 connected to local service processor 116 enabling the inclusion of a date stamp in blade insertion table 152. In the preferred embodiment, RTC 146 is a battery backed persistent clock that remains constantly operational.

When a server blade 100 is inserted into a vacant slot 202 of cabinet 201 thereby altering
25 the state of tamper latch 135, the local service processor 116 performs the insertion log update described above and may send an asynchronous alert. Because the insertion log update is

initiated in response to the insertion event itself, the logging mechanism described herein is independent of any configuration of the server blade 100. Thus, even if a blade is inserted and removed from a slot without ever booting the blade, the blade's insertion log will reflect the insertion.

5 It will be apparent to those skilled in the art having the benefit of this disclosure that the present invention contemplates a system and method for logging resource movements in a data processing network. It is understood that the form of the invention shown and described in the detailed description and the drawings are to be taken merely as presently preferred examples. It is intended that the following claims be interpreted broadly to embrace all the variations of the
10 preferred embodiments disclosed.

IBM CORPORATION
ARMONK, NEW YORK 10506
U.S.A.